AD_____

Award Number:
W81XWH-08-1-0585

TITLE:
Advanced Patient Data Replication and Recovery

PRINCIPAL INVESTIGATOR:
Perez, David
Hendrian, Andrew

CONTRACTING ORGANIZATION:
Eisenhower Medical Center

Rancho Mirage, CA 92270

REPORT DATE:
October 2010

TYPE OF REPORT:
Annual

PREPARED FOR: U.S. Army Medical Research and Materiel Command
Fort Detrick, Maryland 21702-5012

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
01-10-2010	Annual Report	8 SEP 2009 - 7 SEP 2010

4. TITLE AND SUBTITLE
Advanced Patient Data Replication and Recovery

5a. CONTRACT NUMBER

5b. GRANT NUMBER
W81XWH-08-1-0585

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Perez, David

Hendrian, Andrew E.

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Eisenhower Medical Center

Rancho Mirage, CA 92270

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
U.S. Army Medical Research and Materiel Command
Fort Detrick, MD 21702-5012

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

1. The move to electronic medical records (EMR) necessitates that clinical data protection and recovery best practices support extremely low RPO and RTO. Only near-real time and synchronized backups to offsite, disk based, storage technologies support sub-four-hour RPO and RTO while still protecting the data from local loss.

2. Objectives include:
   A. Lower the risk of clinical patient data loss to clinical staff
   B. Support clinicians dependence on EMR data by making it less prone to loss
   C. An "IT healthcare best practice" will be redefined for off-site real-time data replication of electronic medical records which will lower RPO and RTO to less than the current levels of 24 and 48 hours.
   D. Target objectives for RPO and RTO will be 1 to 4 hours.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			USAMRMC
U	U	U	UU	22	**19b. TELEPHONE NUMBER** *(include area code)*

Table of Contents

Introduction

Advanced Patient Data Protection (APDAPT)

Eisenhower Medical Center (EMC) is attempting to lower the risk of losing patient data, as well as the risk incurred by lengthy recovery processes in the case of a data loss, by making available, in near real-time, a duplicate electronic medical record which includes radiological images. EMC has made a multi-million dollar investment in the digitization of patient data; the Electronic Medical Record (EMR). Moving from a process firmly entrenched in the use of paper forms, verbal authorizations, and hand written notes, EMC has digitized the creation, storage and retrieval of the patient chart or EMR. The EMR is comprised of patient vital signs, nurse notes, medications administered, doctors' orders, dietary and radiology orders, radiological studies and results, lab orders and results as well as transcriptions, etc. Undertaking the change to an EMR has required a massive departure from decades-old processes that were intrinsically tied to paper records and manual procedures.

This change necessitates that clinical data protection and recovery best practices support extremely low recovery point objectives (RPO) and recovery time objectives (RTO.) Recovery Point Objective describes the acceptable amount of data loss measured in time.

The Recovery Point Objective (RPO) is the point in time to which you must recover lost data as defined by your organization. This is generally a definition of what an organization determines is an "acceptable loss" in a disaster situation. Traditional backup strategies for many hospitals have focused on business or financial data protection. These strategies considered a 24 hour backup window an acceptable risk. Backups would be created every 24 hours. If a data loss occurred 23 hours and 59 minutes later, there would be no back up for that previous time period back to the previous back up. This represents "lost" data, and the lost data would be recovered from other sources, such as printed reports, statements, etc. This recovery of a "days" worth of data was deemed acceptable. When considering other data types, such as clinical data, that risk is no longer acceptable. For EMC the RPO has been 24 hours. Based on this RPO the data must be restored to within 24 hours of the disaster. All data from the point of the disaster to 24 hours later will have to be manually recovered through other means. In healthcare, with records no longer paper-based, this could prove to be impossibility with some clinical data sets, like verbal pharmacy and lab orders.

The Recovery Time Objective (RTO) is the duration of time and a service level within which a business process must be restored after a disruption in order to avoid unacceptable consequences associated with a break in business continuity. The RTO includes the time for trying to fix the problem without a recovery, the recovery itself, tests, as well as communication to those who use the systems affected. This time frame is usually an objective or goal for an organization, not a mandate. Strategy is often selected that will not meet the RTO. EMC's strategy will be to find a solution that will meet the objective.

While tape based backup have been traditionally the medium of choice, the time it takes to recall off-site tapes and restore from tape can add dozens of hours to a recovery. As such, EMC has decided to investigate a disk-based storage technology that can support sub four-hour RPO and RTO. Moving a fault-tolerant copy of the data offsite using such a technology will protect the data from local loss.

EMC endeavors to achieve the following in regards to clinical data:

1. Lower the risk that EMC will be unable to access patient data from the EMR due to data loss.
2. Reduce the risk of loss of PACS data elements in the event of disaster to EMC's local databases.
3. Increase the availability of the EMR data by lowering the figures for RPO and RTO from 24 and 48+ hours respectively. Target objectives for RPO and RTO will be 1 to 4 hours.
4. Re-define an "IT healthcare best practice" for the protection and recoverability of electronic patient data through the utilization of an off-site, real-time, replication of electronic medical records which will lower RPO and RTO to less than the current levels of 24 and 48 hours (respectively) in the event of data loss within the EMR.

Body and Key Milestones:

Eisenhower Medical Center's approach to reducing risk to the electronic medical record by reducing the time needed to recover from a data loss is to replicate our clinical data to an asynchronous disk-based technology outside of the seismic disaster zone of southern California. As described in the following pages, we have made significant progress in this project. Remote site selection has been completed. Equipment has been shipped and setup at the new location has been completed.

1. Remote site selection visits have been completed and EMC has selected a final location. Data center co-location companies that had multiple sites in multiple states were deemed in the best interest of the organization. Although this criterion was not part of the official selection matrix, it has since been considered. The following vendors were considered:
   - Via-West
   - Core Link Data Centers Inc.
   - SunGuard Data Protection Services
   - MarquisNet Co-Location Inc.

   MarquisNet has been selected as our co-location service provider. Eisenhower Medical Center will utilize their Las Vegas site for our remote site archive location. With locations in Riverside, CA, Phoenix, AZ, and Las Vegas, NV, Eisenhower Medical Center chose MarquisNet's Las Vegas location due to the degree to which this site matched our selection criteria.

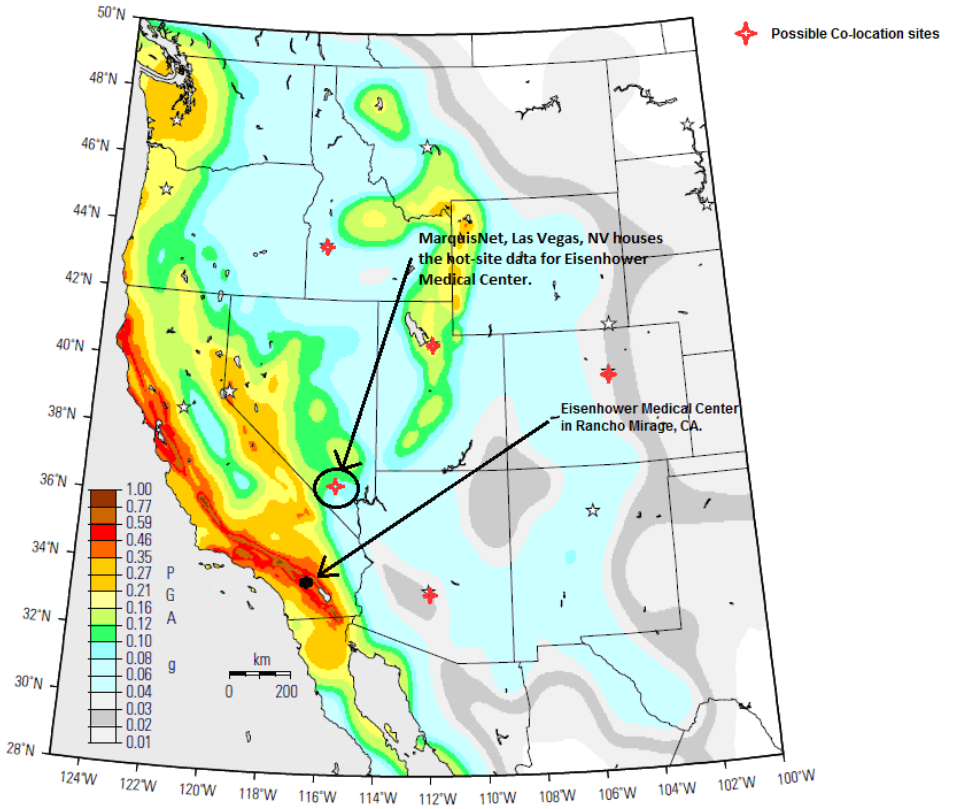

Located at 7185 Pollock Drive, Las Vegas, NV and within 1 mile of the McCarron International Airport and with direct flights from both Palm Springs and Ontario airports, this site is optimal for and meets all identified criterion.

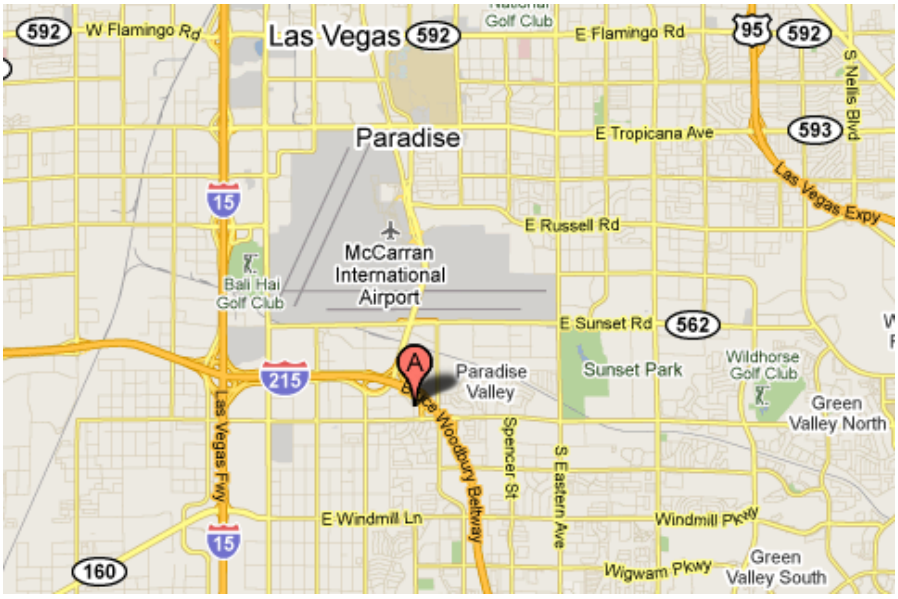

Additionally, MarquisNet's location within Las Vegas was more ideal than that of the other Las Vegas competitors; ViaWest and Core Link Data Centers. Security in MarquisNet building appears by all indications to be more robust and the immediate area surrounding the site also is more appealing to business. A site visit was planned and completed on December 2009. Contract negotiations were completed in March, 2010.

2. The Centera data repository was moved from the Eisenhower location in Rancho Mirage, California to the MarquisNet co-location in Las Vegas, Nevada, arriving March 25, 2010. Additionally, networking hardware was also purchased as well as data-communication links back to EMC. Additionally, 2 extra network racks were also relocated to La Vegas site. These racks were not purchased as part of the grant, but were unused from a previous project so no grant expense was needed for the extra racks.

3. Network connectivity established - EMC staff arrived at the Las Vegas site on March 26 to setup the secure (encrypted) point-to-point network connection back to Eisenhower's data center in California. Additionally, the Centera was attached to the network and replication would be started By March 27, the connection between the EMC campus and Las Vegas data center had been configured. Initial setup was completed the previous day; however, the link would fail repeatedly when replication of data would hit 500Mbps. It would then fail for each 500Mbps of data transmitted. The following "bug" scenario was found on Cisco's troubleshooting site:

   "Highly loaded IPsec tunnel can get stuck after transmitting about ~5500 Mbytes of traffic. When this happens, ""encaps" counters on the affected tunnel get frozen and stop increasing, or increase very slowly (1-2 encaps per second), even if much higher traffic rate is present. "Decaps" counters are still increasing well in accordance to traffic flow. Issue is recovered by "clear crypto IPsec ssl" Cisco IOS command, but occurs again after next 500 Mb. Under-utilized tunnels on the same ASA are not affected. After the failure, send/recv error counters on IPsec SA may start to increase. Also, DPD may appear in ASSA logs after the failure, as there is no traffic in one direction anymore."

   EMC staff upgraded the Cisco ASA firewall IOS (internetworking operating system) from 8.0.4.12 to 8.8.0.5 as recommended by Cisco tech support. We also re-configured the rate limit to pass 25Mbps. Initial configuration was 15Mbps. The IOS upgrade worked and on the 27th we had our stable connection. Catch--up data replication started immediately afterwards. The screen shot bellow illustrates the inbound data rates:

4. Catch-up replication completed – By Monday, March 29 our replication had continued unabated from March 27, averaging nearly 330Mbps throughput. This rapid replication allowed us to "purchase" a smaller bandwidth connection. Our setup connection had an expanded bandwidth or burst data rate of up to 100MMbps. We opted to fully utilize this capacity for the first week of operations. With replication complete well inside of the 7 days, we were able to "dial" the bandwidth back down to the 10Mbps range to reduce

the cost of the data link back to EMC. EMC2 Centera replication is the process whereby an EMC2 Centera cluster is automatically copies new content to another EMC2 Centera (defined in this project as DR Centera.) As an EMC2 Centera cluster acquires new content from a local application, the replication mechanism ensures that this new content is automatically and transparently transferred across a WAN to a designated EMC2 Centera cluster presumably in another location. Eisenhower Medical Center will employ the most basic of EMC2 replication technologies, called unidirectional replication. The application (PACS and HPF) writes data to the local EMC2 cluster and that data is automatically replicated to the DR cluster. Unidirectional replication provides disaster recovery (DR) capabilities where applications may write to a single EMC2 Centera cluster and automatically create and online copy of the data at Eisenhower's remote site in Las Vegas. In case of a disaster or when the primary EMC2 Centera becomes unavailable, the application may failover to the replica DR cluster. Automatic read failover is a feature of the EMC2 Centera and is enabled by default.

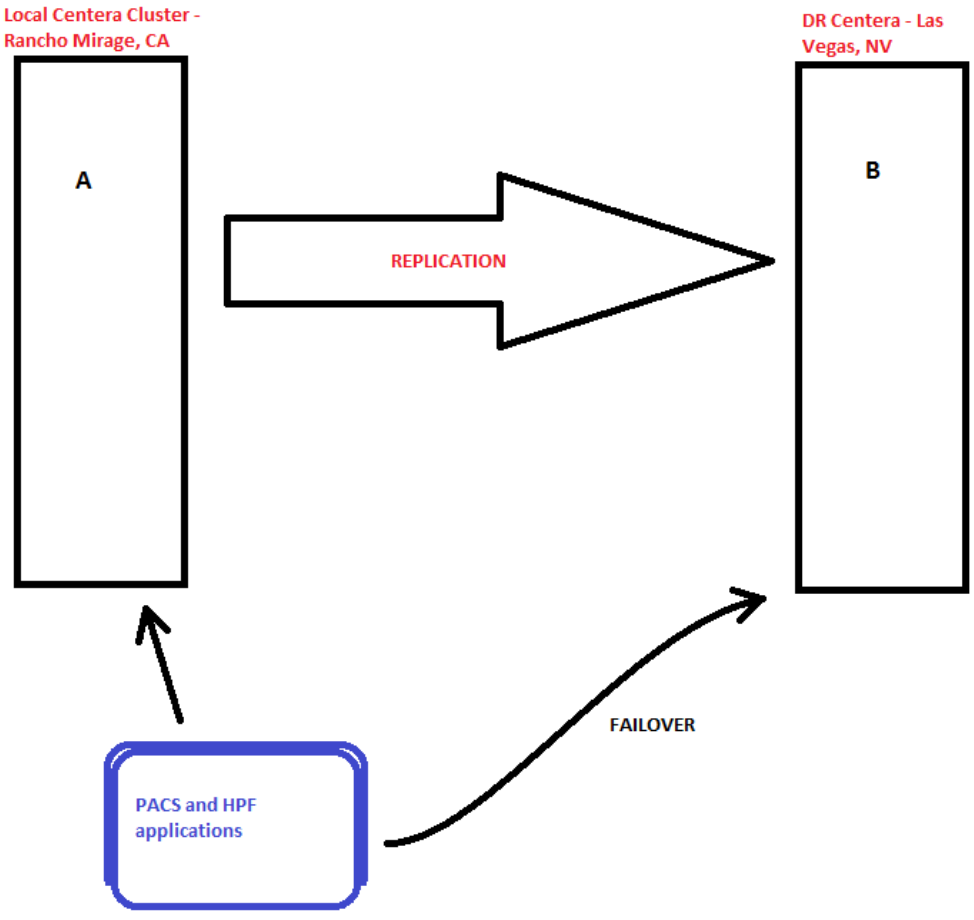

Replication failover

Data written to a DR cluster during a disaster needs to be restored to the original cluster once it is back online. If cluster A was lost during the disaster, all lost data needs to be restored from the DR cluster using the EMC2 Centera

restore feature. As a note of interest, by August 2010, replication notifications (part of the Centera systems) began signaling that replication was out of sync. After a brief investigation, it was determined that the amount of data being replicated had grown and had hit our 10Mbps rate. We reopened the variable link rate to 20Mbps to let the synchronization catch up again. Once in sync, EMC opted to leave the data transmission rate at 20Mbps. Now that the DR target system has been moved to the colocation site in Las Vegas, project achievements (as listed in the Introduction) 1 and 2 have been achieved: 1) Lower the risk that EMC will be unable to access patient data from the EMR due to data loss. 2) Reduce the risk of loss of PACS data elements in the event of disaster to EMC's local databases.

5. Data Failover Procedures – Procedures to create the failover scenario and test the data accessibility are being prepared. These scenarios are prepared in cooperation with technical staff from McKesson Clinical Information Systems support as well as EMC2 technical staff familiar with the Centera storage platform and the

6. Clinician Survey Completed – the survey instrument for gathering of research data has been completed. The survey instrument will be distributed to EMC clinical staff as well clinical staff selected from various attainable mailing lists. The survey instrument is included in Appendix D.

7. IRB Approval – The Request for Human Subjects Research Determination submitted on September 8th, 2010 was prepared and submitted to the Eisenhower Medical Center Institutional Review Board. After an initial interview with the EMC IRB liaison, the IRB viewed the Clinical Survey and the proposed utilization of the survey and found no need for IRB review or approval. "Although the Department of Defense and other medical centers who are converting to an electronic medical record system, the information gathered through this survey is relevant solely to the Eisenhower Medical Center. 45 CFR 46 defines research as a systematic investigation including research development, testing, and evaluation, designed to develop or contribute to generalized knowledge. The procedures described do not qualify under DHHS regulations as research because, although systematic…, the findings of this specific part of the project are not intended to contribute to generalizable knowledge. The IRB review and approval is not required." The full texts of the IRB submission, as well as the IRB Official Determination, are submitted in Appendix C. Unfortunately it is clear in both the IRB Submittal and IRB Determination that the understanding of the use of the information obtained from the survey was incorrect. Throughout the interview with the IRB facilitator it was maintained that the research would result in data generalization in the form of Information Systems Best Practice. Such best practice would be useful to other medical centers using electronic medical records. It is interesting that the IRB Determination surprisingly states that "development of Information Technology Best Practice" (best practice is, by definition, generalizable) was one of the purposes of the research. The contradiction in the IRB Determination demonstrates that

further clarification from the project director is needed in the new IRB submission. Part of the problem is due to the misunderstanding of the research survey human targets. The IRB understood the subjects of the survey research to be limited to the clinical staff of the medical center. In practice, not only will EMC staff be surveyed, but a broader target audience will also be surveyed, resulting in generalizable data. A new submittal will be prepared, fully defining the human survey targets as well as the expectation that the results of the survey will be generalizable research data. Again, the full texts of the IRB submission and the IRB Declaration are found in Appendix C.

8. Finally, the grant timeframe has been extended to April 30th, 2011. The extra time is needed for IRB approval, survey instrument distribution and analysis of instrument results, as well as final project documentation. Please see Appendix E for official document showing grant extension approval.

Milestones not accomplished in the originally expected time frame:

- EMC IRB approval – see accomplishment 7 in preceding section.
- US Army IRB approval – will be submitted as soon as EMC IRB approval is finalized.
- Survey instrument not distributed. This will be accomplished after IRB approval.

Next Steps:

EMC will administer a survey to a wide range of clinicians (nurses, doctors, pharmacists, clinical technicians, etc.) The desired objective is to help set the framework for a clinical IT best practice for disaster recovery priorities for clinical application sets. The survey will be submitted to thousands of clinicians across the country as well as EMC staff clinicians, to cast as wide a net as possible in order to gather as much opinion as possible regarding the relative importance of clinical information systems. The outcome will be to have a standardized and generalizable disaster recovery priority for clinical information and documentation systems. Additionally, EMC will begin testing the failover to the replicated EMR at the remote data center site once testing protocols are prepared, tested, refined and finalized. The failover tests will then be timed and retimed to determine the expected archive failover durations.

Reportable Outcomes:

While the overall objective of the research is still in the early phases, we have garnered some valuable lessons. Almost right off the bat, as project manager, I battled resource constraints as competing project demands strained and then drained away resources from this research project. While the competing projects were prioritized as high visibility by the administration of the hospital, skilled resources were, none the less, needed for this research project. The role of technical project manager proved essentially impossible to backfill. As a result, the research portion of the project was stalled. In order to further resist the drain of resources by future projects we have put the Advanced Data Protection project on the organizations official project portfolio so there is visibility by the administration into this project. Since that time, we have had no resource constraints due to competing projects.

One of the most surprising lessons centered on our research objective of categorizing our clinical applications in regards to disaster recovery. Categorization or prioritization of clinical systems will hopefully give us insight into the recovery order of systems in the event of a major outage. Pulling together a detailed survey of clinicians to identify what, to them, is the most critical clinical applications is central to identifying a healthcare IT best practice for protecting electronic medical records and securing patient safety. We expect to gain much valuable data from this survey. However, in a concurrent, yet non-affiliated project of moving our data center to a newly constructed data center, as we prepared to take down systems preparatory to moving them, our clinician community identified systems critical to their clinical success. This listing of critical systems proved to be surprising. While EMR systems were near the top of the list, the systems involved in the medication administration cycle were identified as *most crucial*. These systems included the pharmacy system, the medication administration bar-coding system, the pharmacy automation robot (which fills pharmacy orders) and the drug dispensing kiosks located near the nurse stations in the hospital. These systems have become so central in patient safety initiatives, and nurses have become so accustomed to the automation achieved, that reverting to manual downtime procedures proved to be unsettling to many nurses. Many of these same nurses, only 4 years previous, were dubious of the need for such systems. Manual processes centering on the medication administration and reconciliation process had to be resurrected and refined in order to achieve the safety level EMC was accustomed to before the data center move. Additionally nurses and pharmacists had to be re-familiarized and retrained in these manual and paper processes and procedures.

Conclusions:

Project results as described in the Report Introduction have been partially fulfilled. Items 1 and 2 are now fully achieved. Now that the secondary data repository is secure and operational, EMC patient data for PACS and the EMR are now protected by duplication both onsite in Rancho Mirage as well as at the co-location site in Las Vegas. Recovery time objective (RTO: the duration within which the EMC EMR business process must be restored after a disruption in order to avoid unacceptable consequences associated with a break in EMR continuity) will be determined and quantified along with the recovery point objective (RPO.) Since the data is now synchronous in two locations there does not appear to be any technical reasons at this point why our original objectives of less than 4 hour RPO/RTO cannot be achieved. However, the study is not far enough in actual testing scenarios to begin analysis.

Additionally, the IT Healthcare best practices will be further refined after administration and analysis of the protocol results. These best practices surrounding prioritization of clinical system protections and restoration among the many clinical systems (EMR's, lab, pharmacy, radiology, transcription, and other systems) will enhance the way in which IT in healthcare can further support the mission-critical nature of healthcare.

Appendix A: Disaster Recovery Remote Site Selection Criteria

1. The remote-site must be outside the regional disaster area of southern California as defined by the United States Geologic Survey National Seismic Hazard Survey. "The 2008 National Seismic Hazard Maps represent the 'best available science' (regarding ground movement vectors and probabilities) based on input from scientists and engineers that participated in the update process." (U.S. Geological Survey Open-File Report 2008–1128: 2008 National Seismic Hazard Maps - Peterson, Mark D., et. al. pg. 40.) See Seismic Hazard map below for relative distances from Eisenhower Medical Center's location in Rancho Mirage, CA. Eisenhower Medical Center and possible remote sites have been superimposed over the National Seismic Hazards map. Note that EMC sits in one of the most at risk areas in the western United States as far as the rate of peak ground acceleration probability. Our goal is to move our secondary archive outside of these high risk areas into areas with much lower peak ground acceleration rates. These USGS and SCEC resources have been instrumental in determining seismic safe-zones nearby to Eisenhower Medical Center.

2.

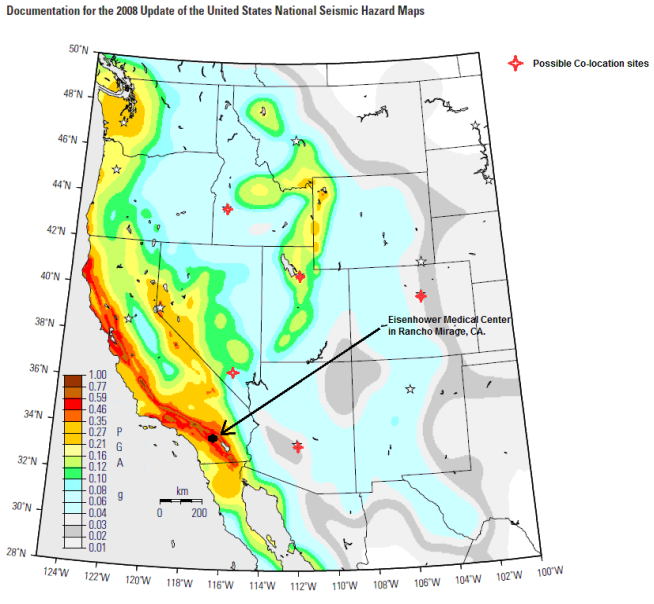

Figure 39. Map of peak ground acceleration (PGA) for 10-percent probability of exceedance in 50 years in the Western United States in standard gravity (g).

3. The remote-site must not be in a coastal area.

4. The remote site city must be in a zone of seismic activity that is less than the seismic activity of the southern California seismic zone. This zone runs the entire length of the state of California and comprises the width (from west coast of Los Angeles basin to the eastern state line due west of L.A.) of southern California.

5. The remote site city must be drivable within 10 hours of EMC via Interstate highway 10, Interstate highway 15, or Interstate highway 8. The naming of these thoroughfares is not a requirement to use them in the event of an emergency. Other roads may be used.

6. The remote site data center must be within a 1-hour drive of a commercial airport in the remote site city that receives traffic from the major commercial freight and passenger airlines.

7. The remote site city should be within 2 hours flight time from either Ontario International Airport in Ontario, CA or Palm Springs International Airport in Palm Springs, CA.

8. The remote site city airport should be accessible via a non-stop flight from Palm Springs airport or Ontario, California airports with Palm Springs airport being the most desired starting point. This is a preference, not a hard requirement.

9. The remote site data center facility should be designed and built to Tier III or greater standards as defined by the Uptime institute. However the site does not need to be certified by the Uptime Institute as Tier III. The major requirements in Tier III standards are: A concurrently maintainable data center has redundant capacity components and multiple independent distribution paths serving the computer equipment. Typically, only one distribution path serves the computer equipment at any time. Additionally, all IT equipment is dual powered and installed properly to be compatible with the topology of the site's architecture. The site would have Tier II or Tier III power utility redundancy and Tier III physical and technical security. A copyrighted Uptime Institute whitepaper titled "Tier Classifications Define Site Infrastructure Performance" has been included in Appendix C. More information regarding the Tier standards can be found there. These principles were heavily relied upon by Eisenhower Medical Center to create these criteria. The Uptime Institute, Inc. is a pioneer in creating and operating knowledge communities for improving uptime effectiveness in data center facilities and information technology organizations. The institute prepares white papers documenting best practices for use by the industry.

10. The site would need to be able to host Eisenhower staff with appropriate workspaces, including telephone and internet service, for several weeks if needed.

11. Hotels must be available within 15 miles or 30 minutes of the selected sites.

12. The remote-site should have multiple data utility Internet Service Providers (ISP) providing service to the hosts.

13. TIA -942 Telecommunications Infrastructure Standards for Data Centers would be required to be followed for electrical grounding and data pathing, fire suppression, networking and cooling for the selected remote site.

14. The ongoing cost of the hosted remote-site must be sustainable by Eisenhower Medical Centers annual operating budget limitations and be approved by the Vice President / CIO of Eisenhower Medical Center.

15. This safety zone for the remote-site has been determined to be at least 200 miles to the north of the Coachella valley or 300 miles east of the valley. Sites west and south were deemed as inappropriate. Westward locations offered no respite from the seismic risks associated with our current location. Southward locations lacked necessary infrastructure to accommodate this projects goals. The most immediate sites are Las Vegas, NV and Phoenix, AZ. Other sites were also reviewed in the cities of Denver, Boise, and Salt Lake City.

Appendix B: Project Budget Expenditures

USAMRAA Pat Data
Replic Recovery 145

Appendix C: IRB Submission to EMC and IRB EMC Declaration

EMC IRB Request - signed.pdf

IRB Appendix C.pdf

Appendix D: Survey Instrument

Questions for DR Clinical survey.pdf

Appendix E: Grant Modification for No Cost Extension

Grant Extension
08-1-0585 P00001.pc